On Routing Instability and End-to-End Path Failures

Nick Feamster

M.I.T. Computer Science and AI Laboratory

David Andersen, Hari Balakrishnan, Frans Kaashoek, *M.I.T.*

Feng Wang, Lixin Gao, *UMass Amherst*

Understanding End-to-End Path Failures

1. Where do end-to-end path failures appear?
2. How long do they last?
3. How do they correlate with BGP instability?
4. How much can path failures be explained by routing?

Data Collection: RON Testbed (~ 30 Hosts)

Active Probes: Detect path failures.

- Pairwise probing; logging detects one-way loss.
- Failure: 2 consecutive lost probes

Traceroutes: Study path IP-level path properties.

- Periodic
- Failure-triggered

BGP Feeds: Detect interdomain routing instability.

- Co-located at 8 measurement hosts

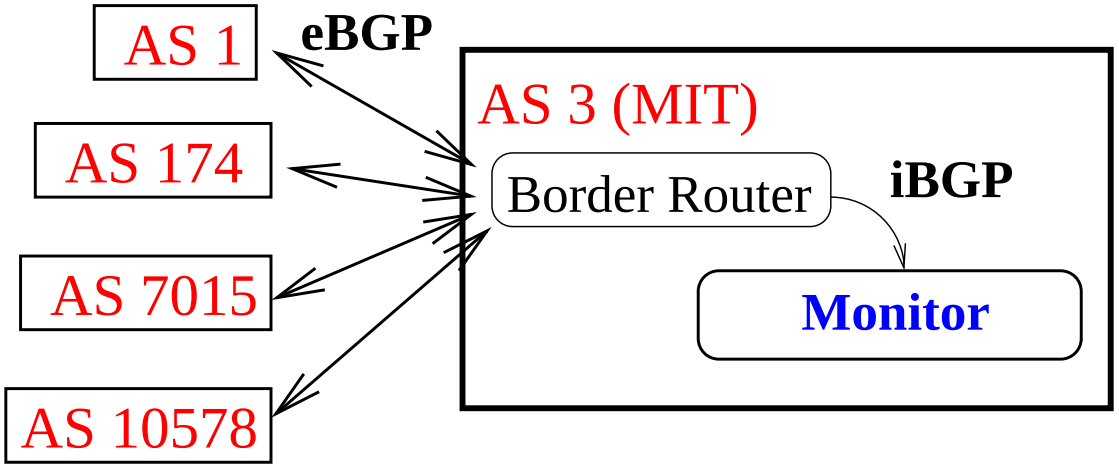

How long do end-to-end path failures last?

90% last less than 15 minutes.

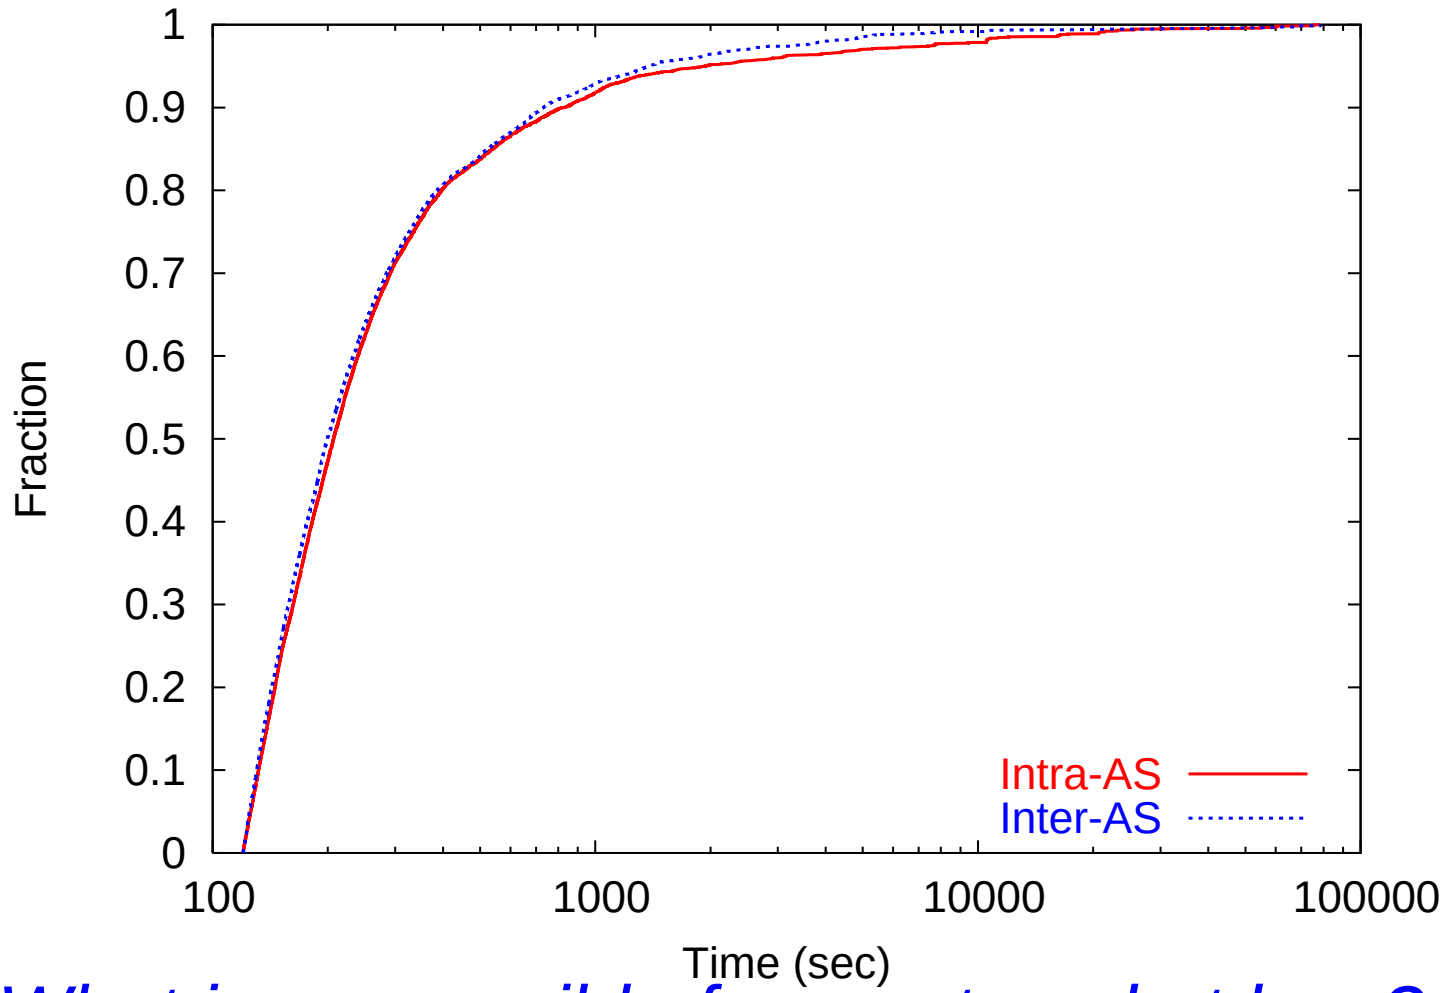

What is responsible for most packet loss?
What causes the long failures?

Routing not responsible for most packet loss

From September 2004 to October 2004:

All Path Failures:

Failure Type	*Number*	*Lost Packets*	*Fraction*
Routing Loops	162	4,991	0.0092
Loop-Free Dynamics	246	24,160	0.0445
Other (e.g., congestion)	331,742	513,862	0.9463

Failures longer than 30 seconds:

Failure Type	*Number*	*Lost Packets*	*Fraction*
Routing Loops	108	4,862	0.0278
Loop-Free Dynamics	150	23,958	0.1372
Other (e.g., congestion)	5,105	145,804	0.8350

Routing dynamics-induced failures last longer

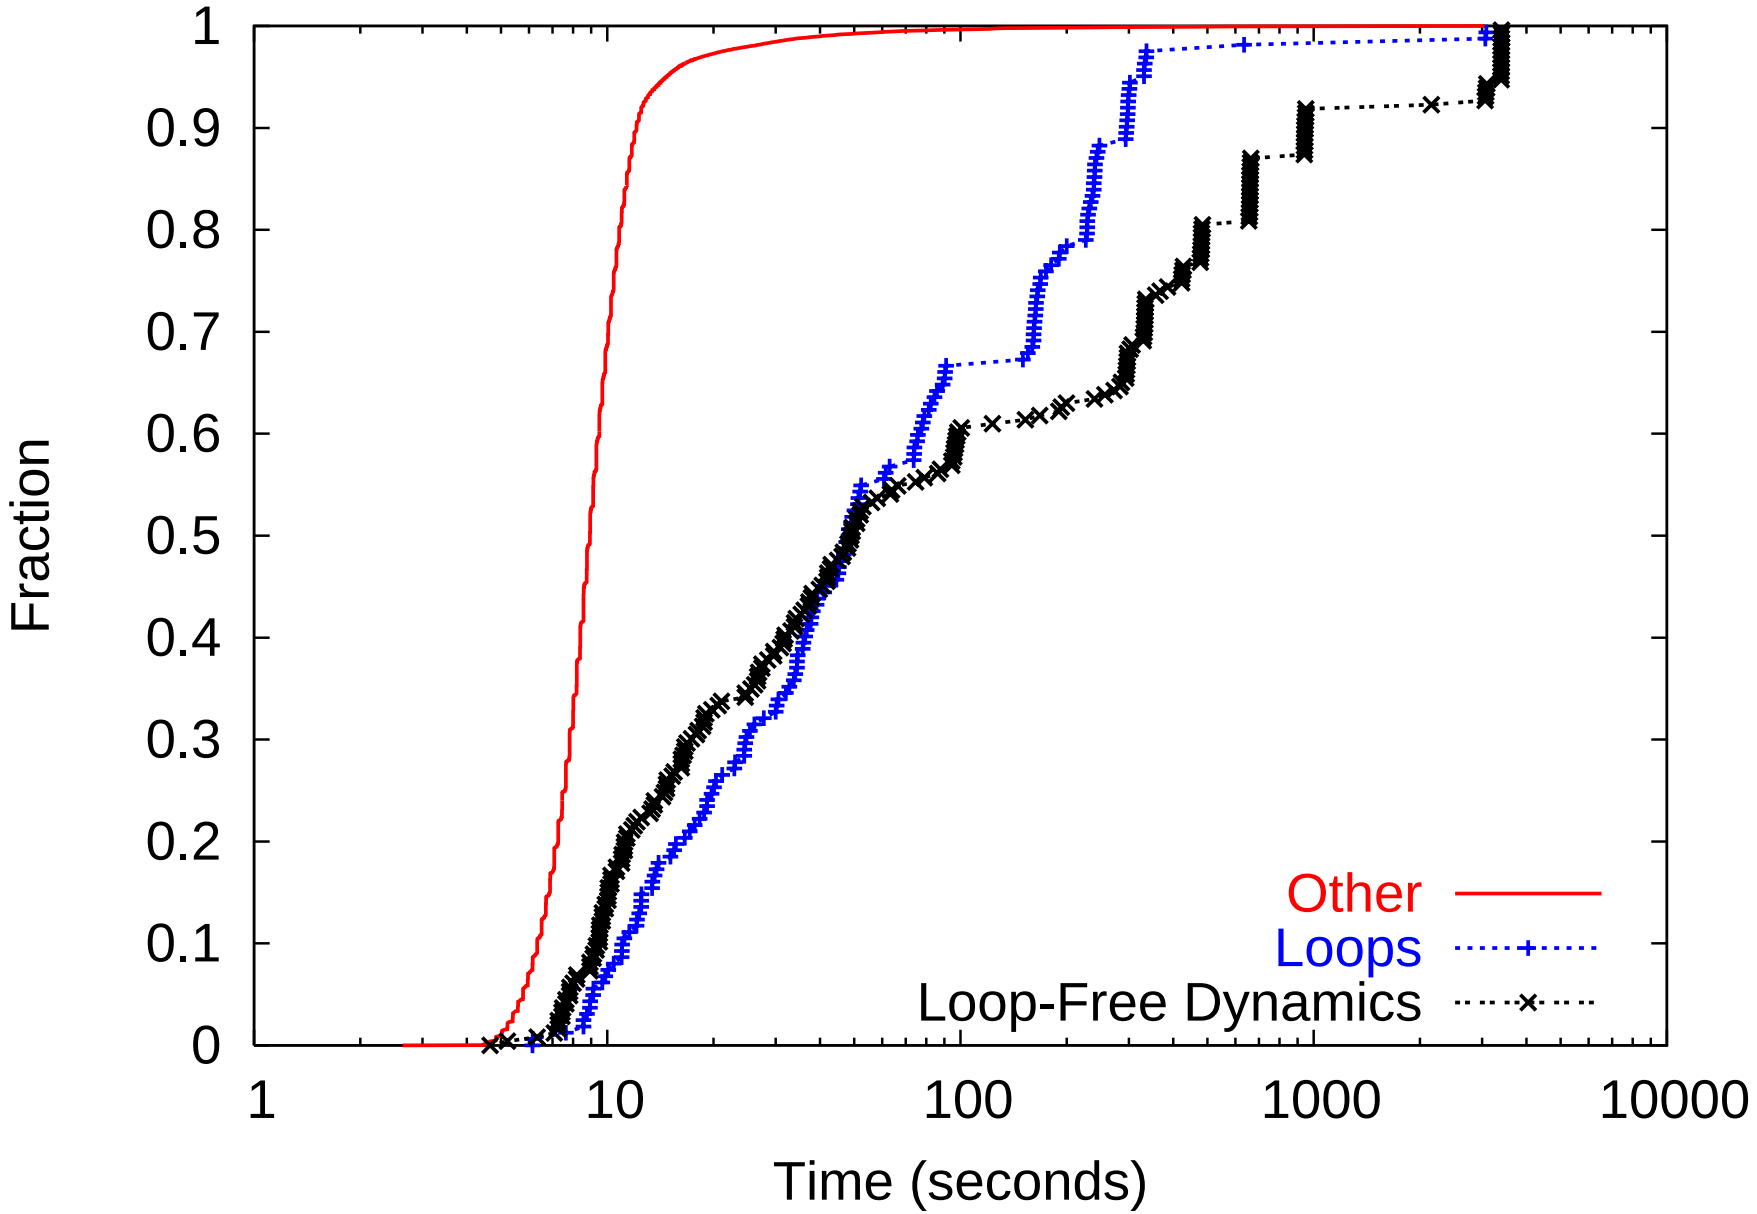

Relating Path Failures and BGP messages

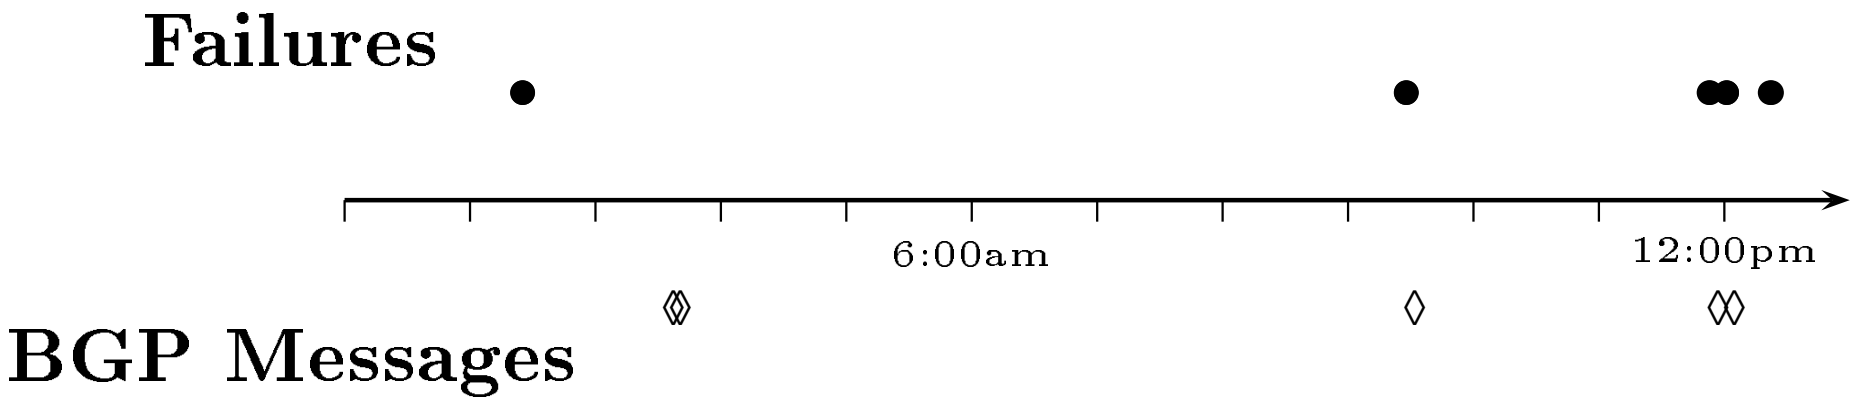

- *Technique 1:* Cross-correlation of time-based signals
- *Technique 2:* Consider a failure and look for BGP (and vice versa)

Do failures correlate with routing instability?

Failures typically occur several minutes before BGP activity.

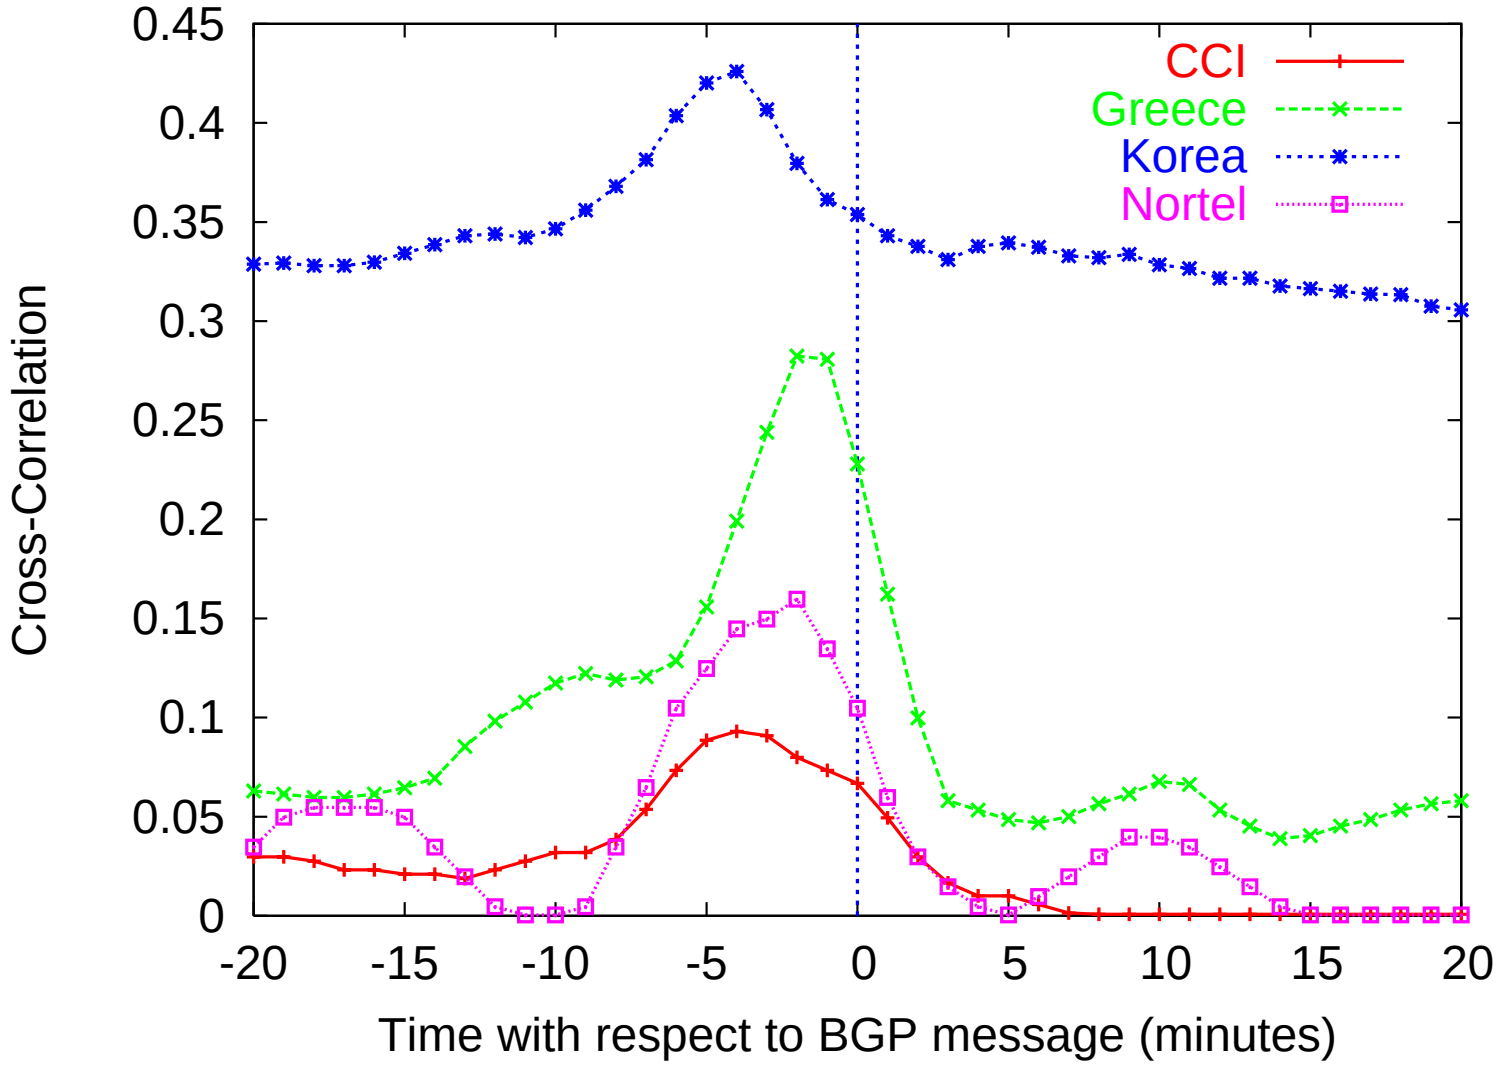

Which failures correlate with instability?

Failures that appear near end hosts are less likely to coincide with BGP instability.

- 60% of failures that appeared at least three hops from an end host coincided with at least one BGP message.
- 22% of failures within one hop of an end host coincided with at least one BGP message.
  - Reachability to an ISP does not imply reachability to customers.
  - These failures are may also be caused by congestion.

Routing dynamics affect independent paths

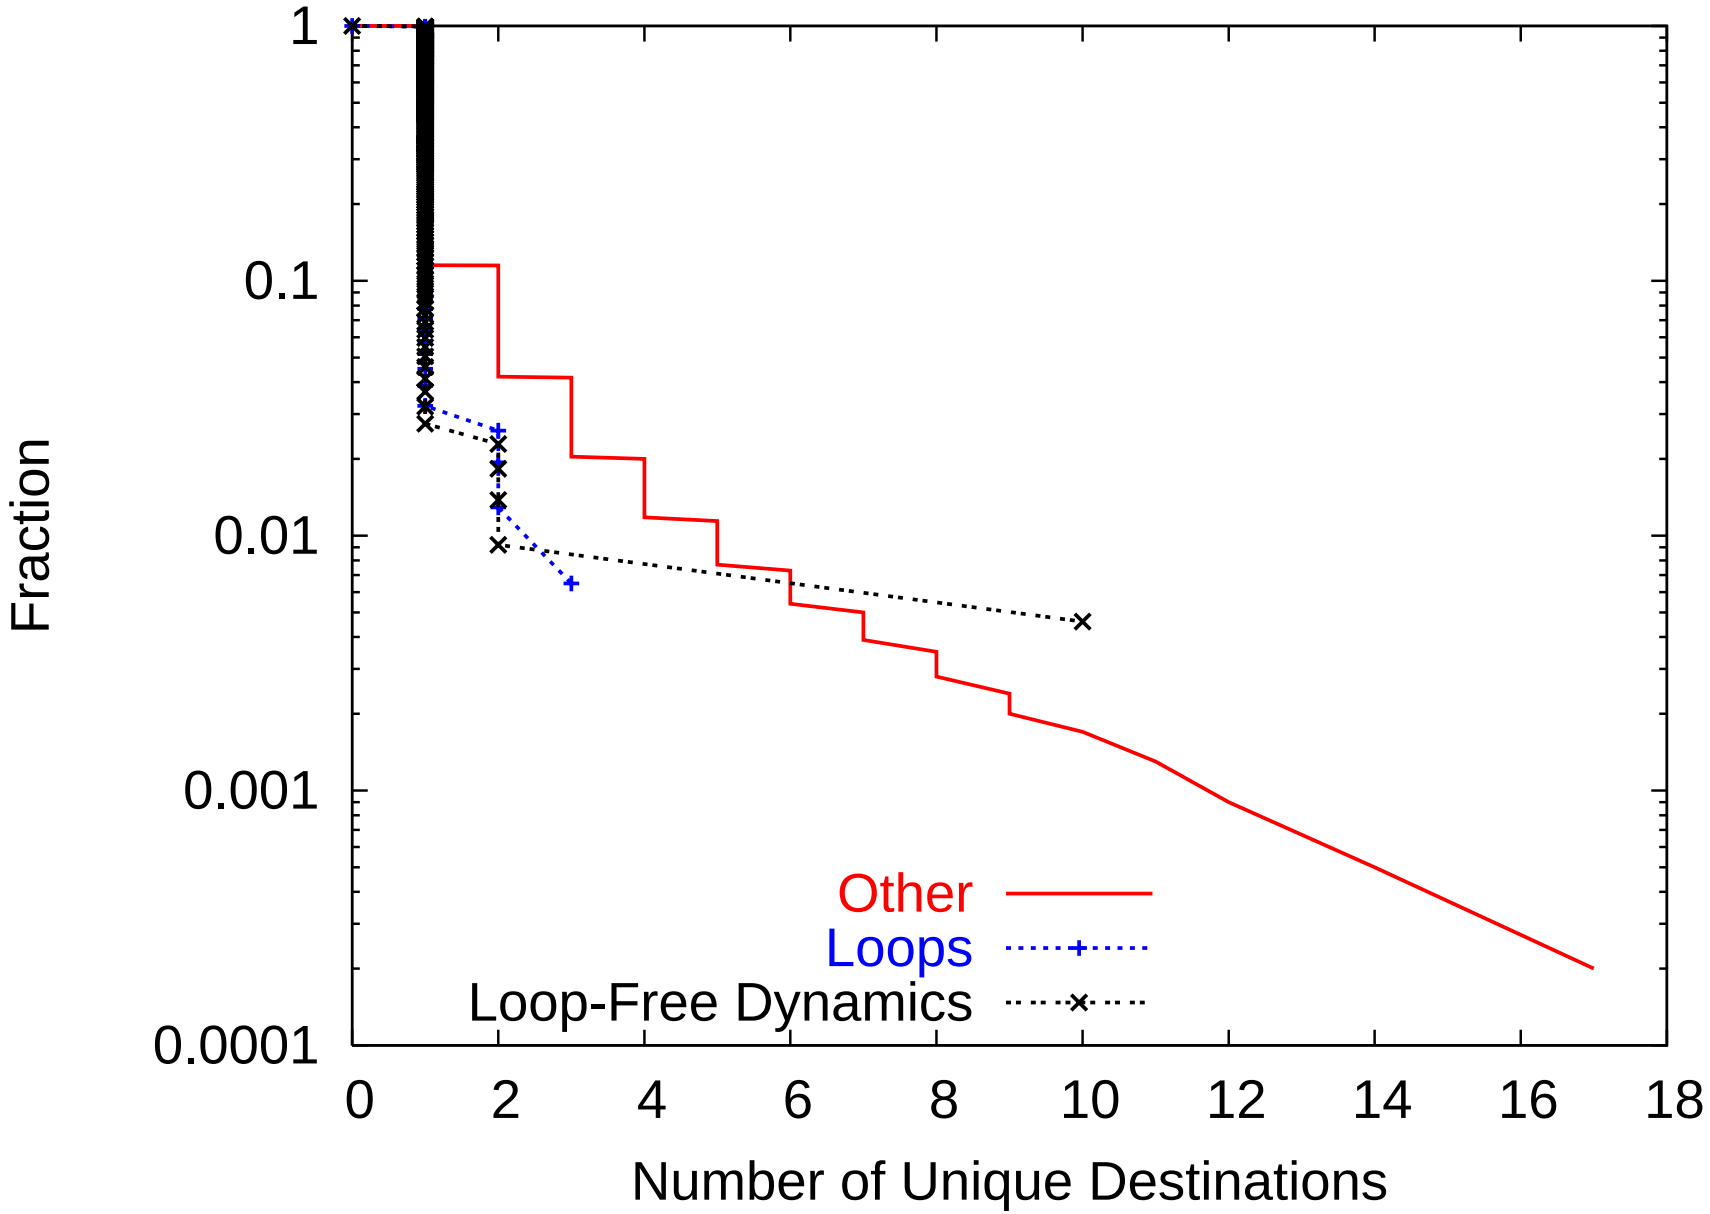

Surprise: BGP messages precede failures!

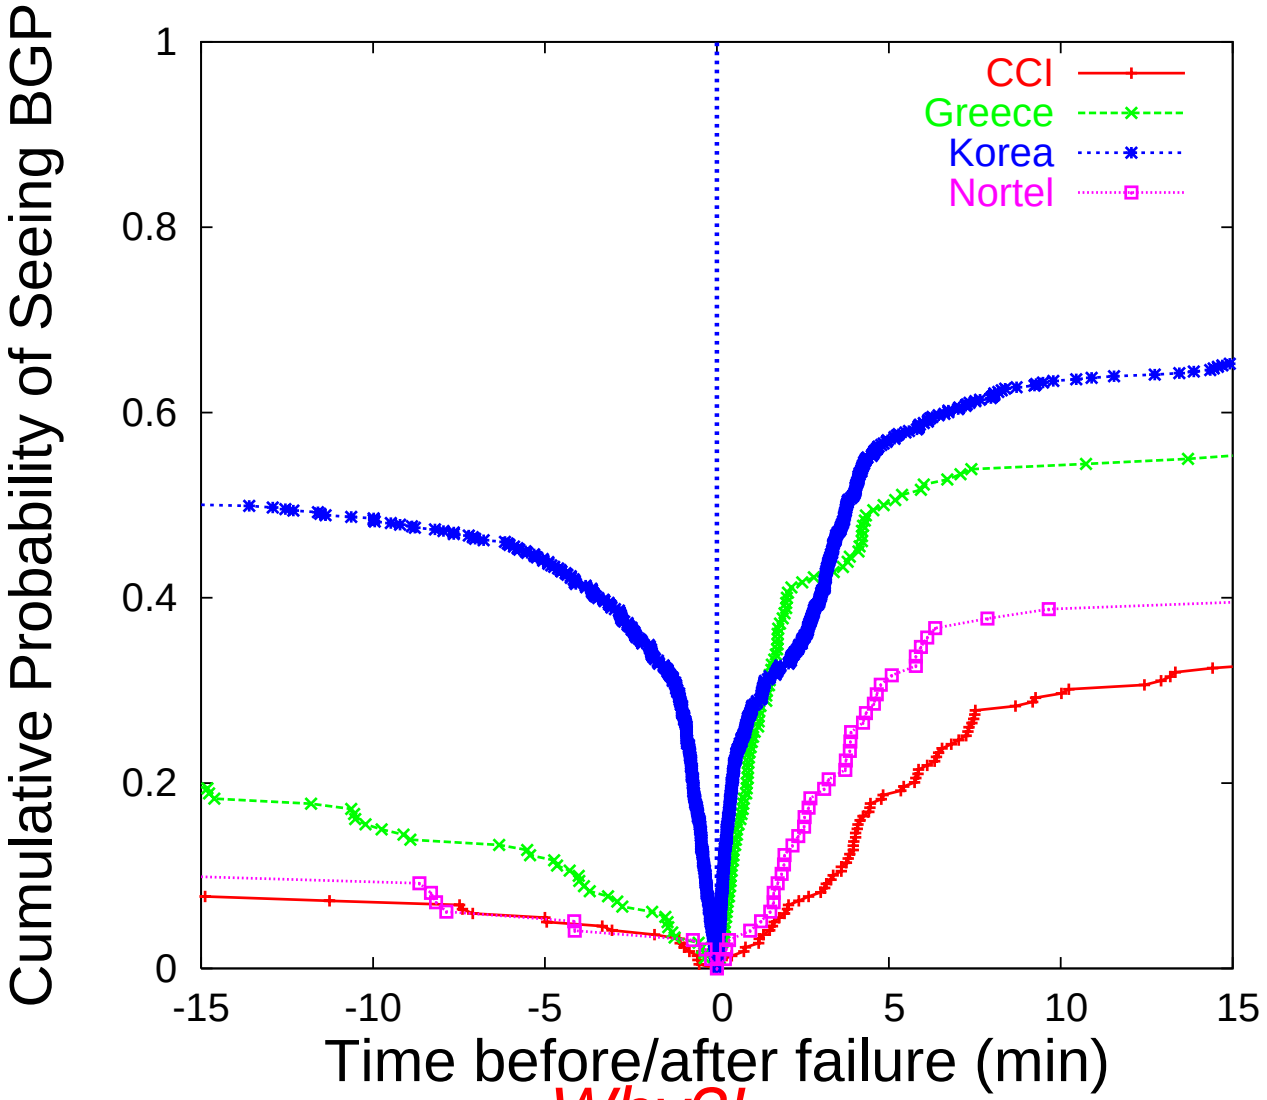

Why?!

Route flap damping, maintenance, misconfiguration, etc.

Can BGP help predict failures?

Effectiveness of predictor depends on path characteristics.

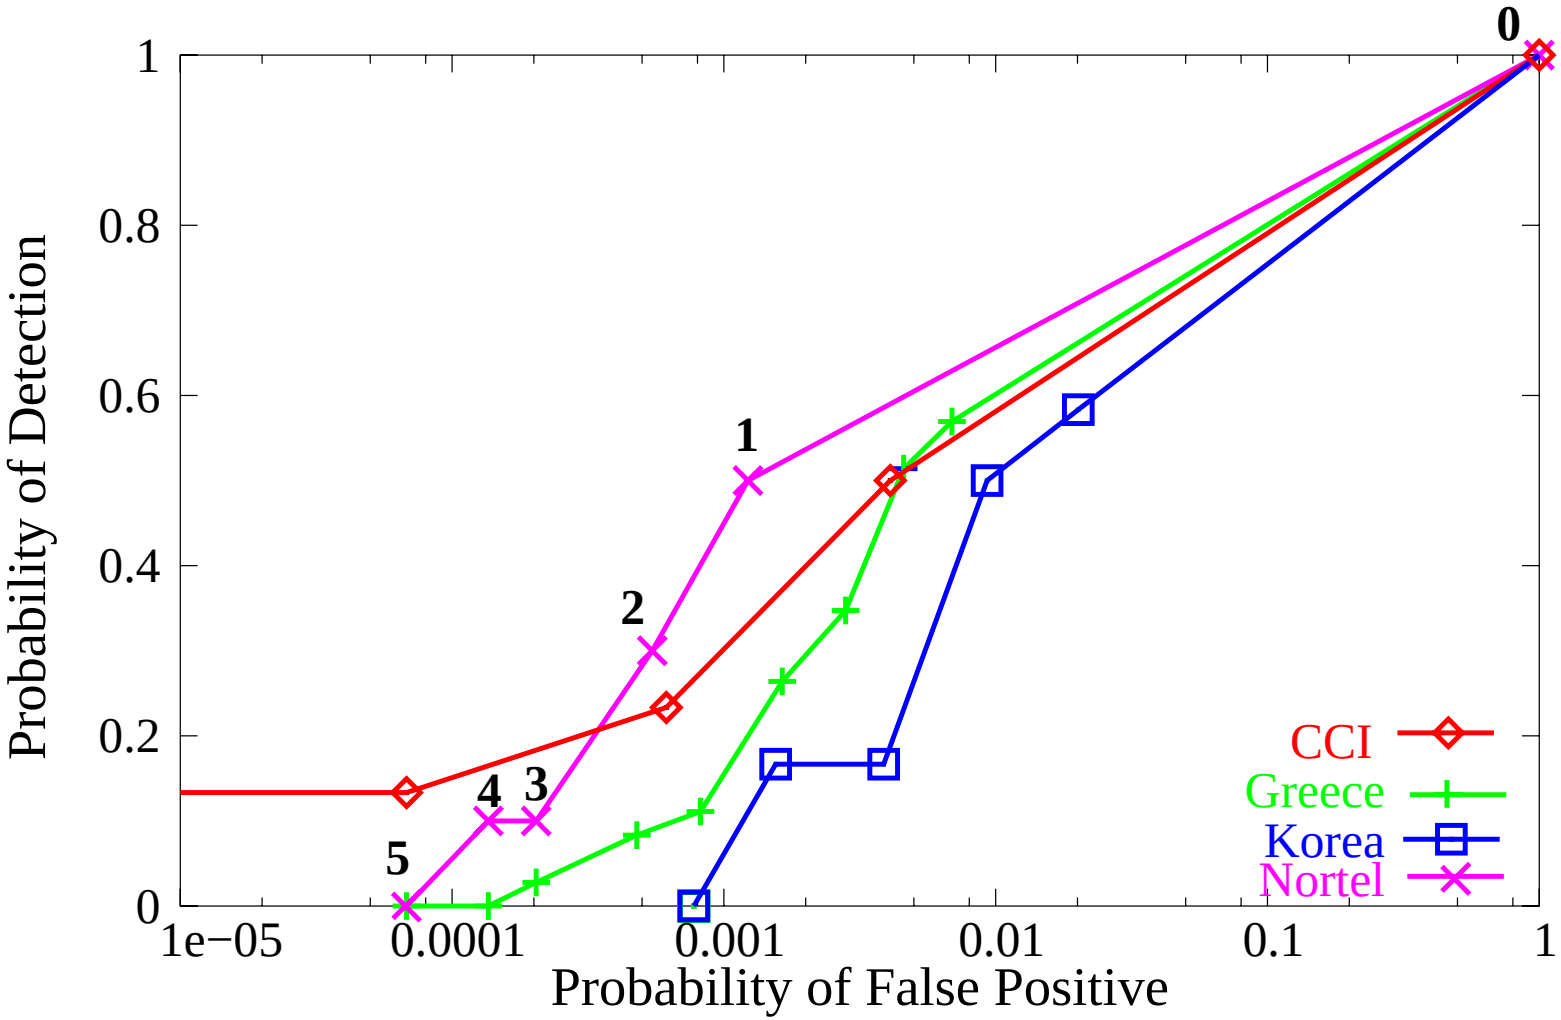

Summary

- *Location*
  - Some links experience many path failures, but many experience some failures.
  - Failures appear more often inside ASes than between them.
  - Congestion-related failures affect more destinations.
- *Duration*
  - 90% of failures last less than 15 minutes
  - 70% of failures last less than 5 minutes
  - Failures caused by routing dynamics last longer
- *Correlation*
  - BGP messages coincide with only half of the failures that reactive routing could potentially avoid.
  - When BGP messages and failures coincide, BGP messages most often follow failures by 4 minutes.
  - BGP sometimes precedes failures.

Failures and delay fluctuations